5 Software Technologies

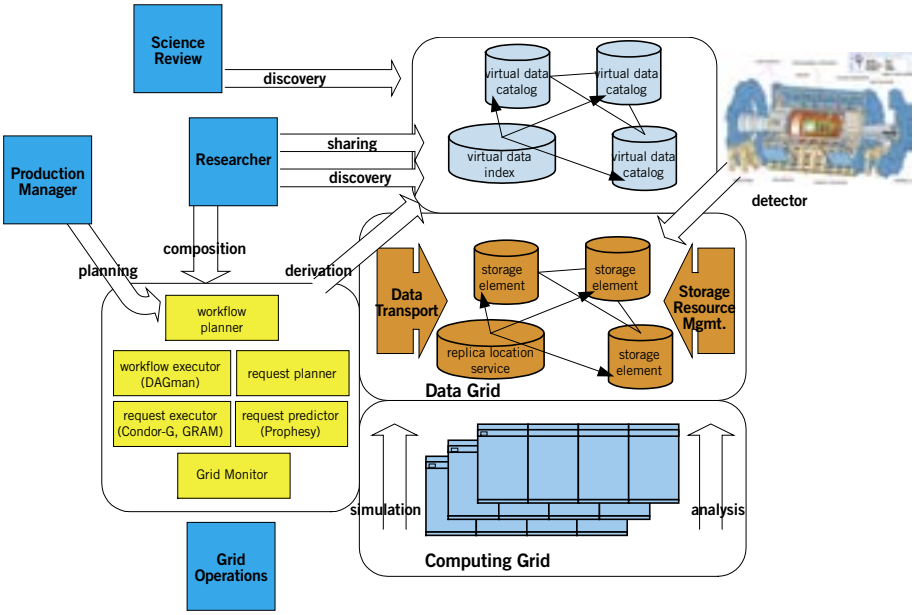

Vision

The software technologies developed and deployed by the Trillium projects focus on achieving scientific productivity in the manipulation and processing of large volumes of data. At the highest level of the toolkit, virtual data systems enable scientists working together throughout the world to discover, understand, manipulate, transform, track, and communicate their data among the members of large, distributed scientific collaborations. At lower levels, these tools turn grids of shared resources into an easy-to-use virtual computing platform, transparently managing and negotiating diverse resource utilization policies.

Components

The central software technologies powering the Trillium Data Grid projects are the Globus Toolkit™, the Condor System, the Storage Resource Manager, the Storage Request Broker, and the Chimera Virtual Data System developed within the research program of the GriPhyN project.

The core of Grid computing platforms is the Globus Toolkit™, consists of four major groups of components for security, information services, resource management, and data management. A suite of tools created by the University of Wisconsin Condor Project provides the advanced scheduling mechanisms of the Data Grid, and works through the GRAM protocol to provide higher-level management of the workflow needed to process requests for virtual datasets.

These components are assembled into the Virtual Data Toolkit, in which all of these components are packaged for ease of installation and for consistent deployment throughout the scientific user community.

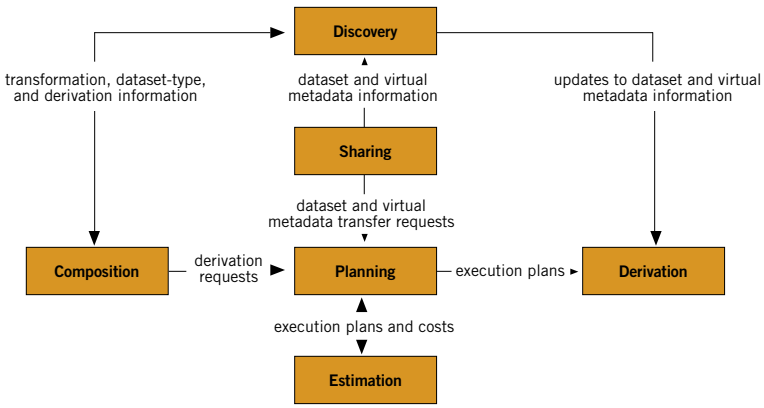

Applications such as the MAGDA distributed data and metadata manager (ATLAS), Sequential Access with MetaData (SAM), MOP Monte Carlo simulation production, Grid Access Portal for Physics Applications (GRAPPA), and Grid Analysis Environment systems, have also been developed to manage application data and computations.

Virtual Data Grid Architecture: A common Virtual Data Grid Architecture, developed jointly by the Trillium Projects, guides the deployment of Virtual Data Toolkit components to build Data Grids that can deliver the envisioned scientific productivity. The architecture guides the assembly of toolkit components into grids that, while highly flexible and tailorable to the needs of specific collaborations, also use globally accepted conventions to enable a high degree of worldwide resource sharing and interoperability, and hence a higher degree of collaboration and knowledge sharing.

...Software Technologies

Grid Security Infrastructure (GSI) implements the protocols that enable secure, trusted, and mutually authenticated communications between multiple grid resource providers over secure sockets using the public key infrastructure. Efforts by PPDG have ensured that public key certificates are accepted in a wide variety of international domains.

Grid Information Services (implemented by MDS, the Monitoring and Discovery Service) provide a scalable information service, which enable grid clients to discover grid computing, storage, and networking resources, and to maintain track of their status.

Grid Resource Allocation and Management services (GRAM) provide a uniform interface to the wide variety of computation resources schedulers. Data management components of the toolkit provide ubiquitous, secure, high-speed parallel data transport through the GridFTP service; a Reliable File Transfer service (RFT) built on top of GridFTP; and the Replica Location Service - a robust, scalable information service to track the location of dataset replicas created around the grid to reduce access latency.

Condor provides a flexible high-throughput scheduling mechanism to manage the compute resources of a Grid site. Its ClassAd service allows jobs to specify their resource requirements and match them to resources at remote grid sites. Condor-G provides a queuing job submission client, which serves as a gateway into the Grid.

DAGman provides the workflow meta-scheduling capability, which manages the parallel-distributed execution of large multi-job workflows. A companion DAPman (data placement manager) component, under development, will schedule the use of storage resources and data staging operations.

Storage Management capabilities to augment the facilities above are provided by the Storage Resource Manager family of components (SRM) from the Lawrence Berkeley Laboratory Scientific Data Management group, and by the Storage Request Broker (SRB, from the San Diego Supercomputing Center).

Research Activities are continuing to extend the capabilities of all the components discussed above towards the vision of ubiquitous, transparent, robust, global computational services. Research into these aspects includes distributed data analysis, fault tolerance techniques, monitoring an prediction algorithms, workload characterization and simulation, data replica placement algorithm design and analysis, troubleshooting, and service definition and discovery using web services protocols. Other research is continuing in distributed data access and data base techniques.

Virtual Data facilities, one of the major research focus areas of the GriPhyN project, are envisioned to become the major user interface to the Data Grid, and to enhance the productivity of science by elevating the manner in which scientists interact with and utilize computing resources. In a virtual data environment, data is managed in units of "datasets," which encapsulate the details of the data's representation and semantic meaning. Tools which process and transform data are encapsulated into transformations with well defined interfaces describing the transformations inputs and output. Work is conducted by composing derivations in which transformations are invoked on datasets to produce new datasets. All derivations are tracked, so that for all datasets in the Grid, scientists can then discover their provenance and can locate both transformations and derivations through powerful information retrieval search mechanisms. This enables new scientists joining a collaboration to discover "what's out there," and to begin, much more easily and rapidly, composing analysis actions to explore the data assets of a collaboration and begin to produce, annotate, and share new results.

Virtual Data Software
- Chimera

Grid Software
- Globus Grid Security
- Globus Resource Management
- Globus Grid Information
- Globus Replica Service
- Configuration and Testing
- Condor Scheduler
- Condor-G User Agent
- DAGman Meta-scheduler
- ClassAds
- Grid Data Mirroring
- PACman Packaging/Deployment

Components Under Development
- Fault Tolerance
- Monitoring
- Policy Management
- Storage Resource Management